Application Notes

hp StorageWorks Data Replication Manager Intersite Link Performance Analyzer

Product Version: ACS Version 8.7P

First Edition (March 2004)

Part Number: AA-RV26A-TE

This document provides instructions for the use of an intersite link performance analyzer tool developed by HP using Microsoft® Excel. This tool will help you determine the worst-case I/O impact over distance in a Data Replication Manager environment.

For the latest version of these application notes and other Data Replication Manager documentation, access the website at http://h18000.www1.hp.com/products/sanworks/drm/index.html. Click the **technical documentation** link and the technical support page is displayed. Click **manuals (guides, supplements, addendums, etc)** for a listing of related documentation.

invent

Data Replication Manager Intersite Link Performance Analyzer Application Notes
First Edition (March 2004)
Part Number: AA-RV26A-TE

About This Document

These application notes cover the following topics:

Intended Audience

This document is intended for customers who are using Data Replication Manager (DRM) and need to estimate the effects of distance on applications that use DRM.

Other Data Replication Manager Documentation

The following documents provide helpful information for running your DRM solution:

- *HP StorageWorks Data Replication Manager HSG80 Version 8.7P Configuration Guide*, part number AA-RPHZF-TE
- *HP StorageWorks Data Replication Manager HSG80 Version 8.7P Failover/Failback Procedures Guide*, part number AA-RPJ0E-TE
- *HP StorageWorks Data Replication Manager HSG80 ACS Release Notes,* part number AA-RPJ2E-TE
- *HP StorageWorks Data Replication Manager HSG80 ACS Version 8.7P Design Guide Reference Guide*, part number AA-RQ78C-TE
- *HP StorageWorks SAN Design Reference Guide*, part number AA-RMPNL-TE

Data Replication Manager Performance Estimation

These application notes provide a methodology and sample results for estimating the effects of distance on application performance when using the HP StorageWorks Data Replication Manager.

A Microsoft Excel tool has been designed by HP to calculate DRM performance based upon cable distances and packet sizes. To interpret and apply the results of the HP StorageWorks Data Replication Manager Intersite Link Performance Analyzer most effectively, you must understand basic DRM I/O. Then you can show the performance impact on the link distance and speed.

DRM Link Variables

For each intersite link, there are two variables that are chosen during the design process:

- The distance between the sites
- The bandwidth of the link

Factors that impact this design are:

- Latency due to the distance between the sites
- I/O write rate (number of writes per second)
- I/O write size (average size in bytes per write)

The performance analyzer tool uses the intersite distance and size of the writes to help in the design of DRM links.

Link Distance

The distance between the sites determines how long it takes for the signal to travel from one site to the other at approximately five microseconds per kilometer per trip. The I/O signal travels the length of the link four times for each I/O transaction. Therefore, the distance will slow down any I/O across the link by adding 20 microseconds per kilometer to the basic zero distance replication.

For long distance DRM, use the driving distance between both sites and multiply that by about 25 to 50 percent. For example, there may be a site where the driving distance is 150 miles and the cable distance is 250 miles. It is the cable distance that needs to be considered and not the driving distance.

Link Bandwidth

The bandwidth determines how long it takes to load the data onto the link. The sum of the start time and the load time equals the time to complete the I/O. The time it takes to load the data onto the link varies depending on the bandwidth of the link and not the length of the link.

Consider how long it takes water to go through a garden hose compared to a fire hose. If each hose is the same length, it will take the same amount of time for the water to come out the other end. Provided each hose can be kept full, the fire hose will transport more gallons per second than the garden hose.

The same physics of distance and pipe size apply to DRM. The time to complete a particular I/O is very dependent on the distance, but is less dependent on the bandwidth. This analogy does not apply in cases of very large I/O and small pipes.

DRM I/O Basic Information

When the DRM initiator controller replicates an I/O to send to the target, a new SCSI I/O request is initiated between the controllers. Because the target storage belongs to the initiator instead of a server, this replicated I/O occurs between controllers and not between the server and controller as in the original I/O.

Given that the I/O is now between the two controllers (an initiator and a target), the I/O follows the standard SCSI over Fibre Channel protocol. A simplified version of the SCSI write protocol consists of the following:

- Initiator to target: *Are you ready?*
- Target to initiator: *Yes I am*.
- Initiator to target: *Here is the data*.
- Target to initiator: *Thanks*.

With this SCSI write protocol, the process makes two round trips across the link to complete the write I/O. Since the data packet is large and the other three transmissions are small, the three small packets are ignored in the throughput calculations.

The link distance determines the time it takes for each of the SCSI protocol packets to travel from one end of the link to the other, and the bandwidth of the link determines how long it takes to load the data onto the cable.

Additional Considerations With Asynchronous Transfer Mode

Additional factors may impact the time it takes an I/O to complete when comparing DRM using Asynchronous Transfer Mode (ATM) to DRM using fiber. Most of these factors are ignored by this performance analyzer. One factor that cannot be ignored is the approximately 300 microseconds it takes to encapsulate the I/O request into the ATM protocol. Since each transaction requires four trips through the link and two conversions per trip, one from Fibre Channel to ATM and one from ATM to Fibre Channel, the conversion by the Open Systems Gateway adds approximately 2.4 milliseconds to the link latency.

Additionally, the ATM link is limited to 155 Mb/s, while direct fiber runs at 1000 Mb/s when constrained by the HSG80 controller. Therefore, while Fibre Channel supports 100 MB/s, a Fibre Channel over ATM link supports approximately 12.6 MB/s (100/8 less an additional 20 percent for ATM overhead). DRM using ATM at T3 or slower speeds is supported but may not be practical from an application perspective.

Link Performance Analyzer

The Link Performance Analyzer is a Microsoft Excel spreadsheet designed to assist you in estimating the DRM intersite link performance. Figure 1 shows the tool when it first opened.

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers): 10 km or 6.2 miles

Enter Average Size of Data Packet (KiloBytes): 2 KB 64 KB max.

	1 Gbps Fiber & 1 GbE IP	OC3 ATM & 100 Mbps IP	T3 ATM	10 Mbps IP
Packet Transmit Time:	0.11 ms	0.27 ms	0.64 ms	2.04 ms
One-Way Inter-Site Latency:	0.05 ms	0.05 ms	0.05 ms	0.05 ms
mSec per I/O:	1.15 ms	2.27 ms	3.64 ms	5.24 ms
I/Os per Second:	868.1 I/O/sec	440.9 I/O/sec	274.7 I/O/sec	190.8 I/O/sec
or	6.25 GB/h	3.17 GB/h	1.98 GB/h	1.37 GB/h
Approximately:	17.36 Mbps	8.82 Mbps	5.49 Mbps	3.82 Mbps
% Bandwidth	1.74%	8.82%	12.21%	38.17%

Figure 1: Link Performance Analyzer

Tool Inputs

Only two inputs are needed for the Link Performance Analyzer:

- Distance between the sites in kilometers. Use 62.5 miles per 100 km (multiply by 8 and then divide by 5).
- The average size of the data packet to be sent in kilobytes. This size is limited to 64 kilobytes, the largest SCSI I/O available. When more than 64 kilobytes of data needs to be written, use multiple writes, 64 kilobytes at a time.

Tool Results

The Link Performance Analyzer displays four columns of results:

- 1 Gbps Fiber & 1 GbE IP
- OC3 ATM & 100 Mbps IP
- T3 ATM (44 Mb/s)
- 10 Mbps IP

Each column performs calculations based on which link is being simulated. The calculations provided are described in the following sections.

Note: Although DRM can use switches running at speeds greater than 1 Gb/s, the links are constrained by the HSG80 controllers that operate at a maximum rate of 1 Gb/s.

Packet Transmit Time

The Packet Transmit Time is the time it takes to load the data onto the pipe. This value is determined by converting the decimal size into a binary size (1 k becomes 1204 bytes) and multiplying by a speed factor of proportionality. This calculation shows the impact of the link bandwidth on the overall performance.

One-Way Intersite Latency

The intersite latency is the time it takes for the I/O to make the four trips across the link. This calculation shows the impact of distance across the link.

mSec per I/O

The mSec per I/O row shows the time in milliseconds for one write I/O to complete. This value is the sum of the first two rows plus any conversion delays, such as the conversion of Fibre Channel to ATM and back. This calculation determines the time it takes to complete one write from the local host to the local controller with synchronous replication to the remote controller.

I/Os per Second

The single stream I/O per second is determined by inverting the mSec per I/O value.

The product of the I/Os per second and the packet size approximates the volume of data that can be pushed through the link over a specified amount of time by a single I/O stream. During remote copy set normalization, the I/O size is 128 blocks of 512 bytes each or 64 kilobytes.

Simplifying Assumptions

The calculations assume:

- There are single I/O streams.
- The application is serial.
- The application issues one I/O at a time.

Most applications, such as Oracle or Microsoft Exchange, can issue multiple I/O streams, so the results of the single I/O stream can be multiplied by the number of parallel streams that can be sustained.

Final Considerations

Read performance across the intersite link takes at least as long as a write because the data may not be in cache. A write is completed when the data is in cache, not when it is actually written to disk. Therefore, cross link read performance can be estimated using the same data as the synchronous writes.

To estimate how many I/O streams are required to keep the pipe filled, divide the milliseconds per I/O by the packet transmit time as the first estimate. An upper limit is imposed by the bandwidth of the pipe or the number of outstanding I/Os available from the controller and host bus adapter.

Performance Analyzer Examples

This section provides examples of results generated by the HP StorageWorks Data Replication Manager Intersite Link Performance Analyzer using various inputs. The following is a list of sample data provided:

- Figure 2: 0 km and 2 Kbytes
- Figure 3: 0 km and 4 Kbytes
- Figure 4: 0 km and 8 Kbytes
- Figure 5: 10 km and 2 Kbytes
- Figure 6: 10 km and 4 Kbytes
- Figure 7: 10 km and 8 Kbytes
- Figure 8: 100 km and 2 Kbytes
- Figure 9: 100 km and 4 Kbytes
- Figure 10: 100 km and 8 Kbytes

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	0	km or	0.0	miles
Enter Average Size of Data Packet (KiloBytes):	2	KB	64 KB max.	

	1 Gbps Fiber & 1 GbE IP		OC3 ATM & 100 Mbps IP		T3 ATM		10 Mbps IP	
Packet Transmit Time:	0.11	ms	0.27	ms	0.64	ms	2.04	ms
One-Way Inter-Site Latency:	0.00	ms	0.00	ms	0.00	ms	0.00	ms
mSec per I/O:	0.95	ms	2.07	ms	3.44	ms	5.04	ms
I/Os per Second:	1050.4	I/O/sec	483.6	I/O/sec	290.7	I/O/sec	198.4	I/O/sec
or	7.56	GB/h	3.48	GB/h	2.09	GB/h	1.43	GB/h
Approximately:	21.01	Mbps	9.67	Mbps	5.81	Mbps	3.97	Mbps
% Bandwidth	2.10%		9.67%		12.92%		39.68%	

Figure 2: Example using 0 km and 2 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	0	km or	0.0 miles
Enter Average Size of Data Packet (KiloBytes):	4	KB	64 KB max.

	1 Gbps Fiber & 1 GbE IP		OC3 ATM & 100 Mbps IP		T3 ATM		10 Mbps IP	
Packet Transmit Time:	0.22	ms	0.54	ms	1.28	ms	4.08	ms
One-Way Inter-Site Latency:	0.00	ms	0.00	ms	0.00	ms	0.00	ms
mSec per I/O:	1.06	ms	2.34	ms	4.08	ms	7.08	ms
I/Os per Second:	939.8	I/O/sec	428.1	I/O/sec	245.1	I/O/sec	141.2	I/O/sec
or	13.53	GB/h	6.16	GB/h	3.53	GB/h	2.03	GB/h
Approximately:	37.59	Mbps	17.12	Mbps	9.80	Mbps	5.65	Mbps
% Bandwidth	3.76%		17.12%		21.79%		56.50%	

Figure 3: Example using 0 km and 4 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	0	km or	0.0 miles
Enter Average Size of Data Packet (KiloBytes):	8	KB	64 KB max.

	1 Gbps Fiber & 1 GbE IP		OC3 ATM & 100 Mbps IP		T3 ATM		10 Mbps IP	
Packet Transmit Time:	0.45	ms	1.07	ms	2.56	ms	8.16	ms
One-Way Inter-Site Latency:	0.00	ms	0.00	ms	0.00	ms	0.00	ms
mSec per I/O:	1.29	ms	2.87	ms	5.36	ms	11.16	ms
I/Os per Second:	776.4	I/O/sec	348.2	I/O/sec	186.6	I/O/sec	89.6	I/O/sec
or	22.36	GB/h	10.03	GB/h	5.37	GB/h	2.58	GB/h
Approximately:	62.11	Mbps	27.86	Mbps	14.93	Mbps	7.17	Mbps
% Bandwidth	6.21%		27.86%		33.17%		71.68%	

Figure 4: Example using 0 km and 8 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	10	km or	6.2 miles
Enter Average Size of Data Packet (KiloBytes):	2	KB	64 KB max.

	1 Gbps Fiber & 1 GbE IP	OC3 ATM & 100 Mbps IP	T3 ATM	10 Mbps IP
Packet Transmit Time:	0.11 ms	0.27 ms	0.64 ms	2.04 ms
One-Way Inter-Site Latency:	0.05 ms	0.05 ms	0.05 ms	0.05 ms
mSec per I/O:	1.15 ms	2.27 ms	3.64 ms	5.24 ms
I/Os per Second:	868.1 I/O/sec	440.9 I/O/sec	274.7 I/O/sec	190.8 I/O/sec
or	6.25 GB/h	3.17 GB/h	1.98 GB/h	1.37 GB/h
Approximately:	17.36 Mbps	8.82 Mbps	5.49 Mbps	3.82 Mbps
% Bandwidth	1.74%	8.82%	12.21%	38.17%

Figure 5: Example using 10 km and 2 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	10	km or	6.2 miles
Enter Average Size of Data Packet (KiloBytes):	4	KB	64 KB max.

	1 Gbps Fiber & 1 GbE IP	OC3 ATM & 100 Mbps IP	T3 ATM	10 Mbps IP
Packet Transmit Time:	0.22 ms	0.54 ms	1.28 ms	4.08 ms
One-Way Inter-Site Latency:	0.05 ms	0.05 ms	0.05 ms	0.05 ms
mSec per I/O:	1.26 ms	2.54 ms	4.28 ms	7.28 ms
I/Os per Second:	791.1 I/O/sec	394.3 I/O/sec	233.6 I/O/sec	137.4 I/O/sec
or	11.39 GB/h	5.68 GB/h	3.36 GB/h	1.98 GB/h
Approximately:	31.65 Mbps	15.77 Mbps	9.35 Mbps	5.49 Mbps
% Bandwidth	3.16%	15.77%	20.77%	54.95%

Figure 6: Example using 10 km and 4 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	10	km or	6.2 miles
Enter Average Size of Data Packet (KiloBytes):	8	KB	64 KB max.

	1 Gbps Fiber & 1 GbE IP	OC3 ATM & 100 Mbps IP	T3 ATM	10 Mbps IP
Packet Transmit Time:	0.45 ms	1.07 ms	2.56 ms	8.16 ms
One-Way Inter-Site Latency:	0.05 ms	0.05 ms	0.05 ms	0.05 ms
mSec per I/O:	1.49 ms	3.07 ms	5.56 ms	11.36 ms
I/Os per Second:	672.0 I/O/sec	325.5 I/O/sec	179.9 I/O/sec	88.0 I/O/sec
or	19.35 GB/h	9.38 GB/h	5.18 GB/h	2.54 GB/h
Approximately:	53.76 Mbps	26.04 Mbps	14.39 Mbps	7.04 Mbps
% Bandwidth	5.38%	26.04%	31.97%	70.42%

Figure 7: Example using 10 km and 8 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	100	km or	62.1 miles
Enter Average Size of Data Packet (KiloBytes):	2	KB	64 KB max.

	1 Gbps Fiber & 1 GbE IP	OC3 ATM & 100 Mbps IP	T3 ATM	10 Mbps IP
Packet Transmit Time:	0.11 ms	0.27 ms	0.64 ms	2.04 ms
One-Way Inter-Site Latency:	0.50 ms	0.50 ms	0.50 ms	0.50 ms
mSec per I/O:	2.95 ms	4.07 ms	5.44 ms	7.04 ms
I/Os per Second:	338.8 I/O/sec	245.8 I/O/sec	183.8 I/O/sec	142.0 I/O/sec
or	2.44 GB/h	1.77 GB/h	1.32 GB/h	1.02 GB/h
Approximately:	6.78 Mbps	4.92 Mbps	3.68 Mbps	2.84 Mbps
% Bandwidth	0.68%	4.92%	8.17%	28.41%

Figure 8: Example using 100 km and 2 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	100	km or	62.1	miles
Enter Average Size of Data Packet (KiloBytes):	4	KB	64 KB max.	

	1 Gbps Fiber & 1 GbE IP		OC3 ATM & 100 Mbps IP		T3 ATM		10 Mbps IP	
Packet Transmit Time:	0.22	ms	0.54	ms	1.28	ms	4.08	ms
One-Way Inter-Site Latency:	0.50	ms	0.50	ms	0.50	ms	0.50	ms
mSec per I/O:	3.06	ms	4.34	ms	6.08	ms	9.08	ms
I/Os per Second:	326.4	I/O/sec	230.6	I/O/sec	164.5	I/O/sec	110.1	I/O/sec
or	4.70	GB/h	3.32	GB/h	2.37	GB/h	1.59	GB/h
Approximately:	13.05	Mbps	9.23	Mbps	6.58	Mbps	4.41	Mbps
% Bandwidth	1.31%		9.23%		14.62%		44.05%	

Figure 9: Example using 100 km and 4 KB data packet

Performance Estimator - HP StorageWorks Data Replication Manager
Peak Single Stream I/O for DRM over IP, ATM, or 1 Gbps Fiber/WDM
Estimates for Synchronous Remote Reads or Replication of Local Writes

Enter Cable Distance between sites (kilometers):	100	km or	62.1	miles
Enter Average Size of Data Packet (KiloBytes):	8	KB	64 KB max.	

	1 Gbps Fiber & 1 GbE IP		OC3 ATM & 100 Mbps IP		T3 ATM		10 Mbps IP	
Packet Transmit Time:	0.45	ms	1.07	ms	2.56	ms	8.16	ms
One-Way Inter-Site Latency:	0.50	ms	0.50	ms	0.50	ms	0.50	ms
mSec per I/O:	3.29	ms	4.87	ms	7.36	ms	13.16	ms
I/Os per Second:	304.1	I/O/sec	205.3	I/O/sec	135.9	I/O/sec	76.0	I/O/sec
or	8.76	GB/h	5.91	GB/h	3.91	GB/h	2.19	GB/h
Approximately:	24.33	Mbps	16.42	Mbps	10.87	Mbps	6.08	Mbps
% Bandwidth	2.43%		16.42%		24.15%		60.79%	

Figure 10: Example using 100 km and 8 KB data packet